The Electronic Curator or How to Ride Your CycleGAN

Eyal Gruss
Mahanaim 134 Tel-Aviv University
eyalgruss@gmail.com

Eran Hadas
Mahanaim 134 Tel-Aviv University
ehadas@gmail.com

Abstract

The Electronic Curator examines whether a computer can not only generate art, but also evaluate its quality. [1] The work uses a Generative Adversarial Network (GAN), which constitutes a dialog between two competing neural networks. Here one represents a painter, who turns a human face into a vegetable portrait (fig. 1). The other represents a curator, who evaluates whether the portrait indeed looks like vegetable faces and encourages the painter to improve. The dialog between the competing networks represents the artistic process.

Training is unsupervised based on the cycle-consistent generative adversarial networks (CycleGAN). [2] Thus we require only a set of face images and an unpaired and unrelated small set of vegetable-faces collected from a Google search on the Internet (fig. 2). In order to avoid mode collapse and get diverse and interesting results, we use a modified loss function inspired by DistanceGAN. [3]

In exhibition mode, the painter observes the spectator's face and turns it in real time into a vegetable-face. The curator then grades the outcome. If the outcome is good enough to confuse the curator, a curatic text is generated based on the vegetables and fruits found in the portrait by object detection (fig. 3). In a world of AI art and creative machines, will the art of curation remain reserved for humans?

In the talk, we will review the techniques that helped in training and in inference, as well as those which did not help. Namely, we will discuss data collection and training strategies, modifications to the loss, and inference time normalization.

Fig 1. A vegetable-face generated in real-time in inference.

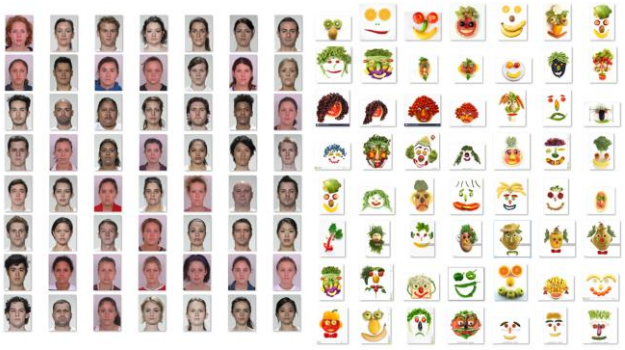

Fig 2. Unpaired samples from the training set in the two domains.

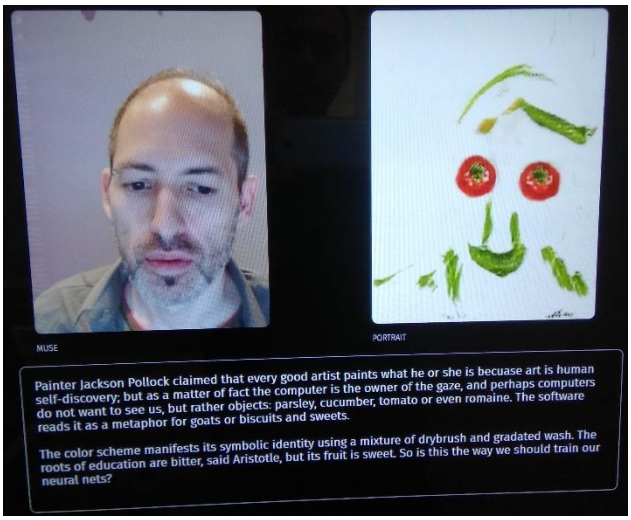

Fig 3. The first author's pretty face, its corresponding vegetable portrait, and the curatic text generated for it. From an exhibition at Heinz Nixdorf MuseumsForum, Paderborn, Germany.

References

1. Project video: youtube.com/watch?v=4sZsx4FpMxg.
2. Zhu, Park, Isola and Efros, Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks, junyanz.github.io/CycleGAN.
3. Benaim and Wolf, One-Sided Unsupervised Domain Mapping, arxiv.org/abs/1706.00826.

Biographies

Eyal Gruss is a machine learning researcher and an artist. He is based in Israel and holds a PhD in physics. His works include poetry, interactive installations and computer-generated art.

Eran Hadas is an Israeli poet, software developer and media artist. Among his collaborative projects are a headset that generates poems from brainwaves (with Gruss), and a documentarian robot that interviews people about the meaning of being human. Hadas was the 2017 Schusterman Artist-in-Residence at Caltech. He teaches in the New Media Program at Tel-Aviv University.

Mahanaim 134 is Gruss and Hadas' tech-art collaboration.